


CONNEXIONS[®]



The Interoperability Report

December 1996

Volume 10, No. 12

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Internet Digital Libraries.....	2
Internet-II Vision.....	8
Internet-II Architecture.....	11
Internet-II Scope.....	14
100Mbps Ethernet Tests.....	16
TERENA.....	25
SNMPng Status Report.....	34
Book Review.....	37
Announcements.....	38

ConneXions is published monthly by Interop Company, a division of SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California, 94404-1138, USA.
Phone: +1 (415) 578-6900
Fax: +1 (415) 525-0194
E-mail: connexions@interop.com

Subscription hotline: 1-800-575-5717
or +1 610-892-1959

Copyright © 1996 by Interop Company.
Quotation with attribution encouraged.
ConneXions—The Interoperability Report
and the *ConneXions* logo are registered
trademarks of Interop Company.

ISSN 0894-5926

From the Editor

As I write the final in-print editorial for *ConneXions*, I am tempted to look back over the last ten years, discuss trends, look at major milestones for the Internet and its technologies, or re-publish a few “classic” articles. But such an approach is probably better left to the new medium of the World-Wide Web where readers can jump both in time and place from article to article. So in this issue, we will look ahead to the future instead.

We begin by discussing the concept of a *Digital Library*. Our collection of back issues will eventually become an example of such a library. The article is by Jack Kessler who serves as an advisor to *ConneXions* as we transition from print to online.

The Internet continues to grow at an amazing rate, giving new users all over the globe access to a growing number of services. For some users, particularly those in research and academia, this popularization has had a serious impact on those who transfer large amounts of data between sites, or who use streaming applications such as packet video for conferencing or collaboration. The concept of a “replacement Internet,” known as *Internet-II*, for the academic and research community has emerged. In three articles we look at some of the details.

Last month we published a tutorial on the two 100 Mbps Ethernet technologies. This month we bring you the results of some tests that compare the performance of these two LAN standards.

The *Trans-European Research and Education Networking Association* (TERENA) was established in 1994 through the merger of the RARE and EARN associations. We asked the Secretary General of TERENA, Karel Vietsch, to give us an overview of TERENA’s activities, its role in the European networking environment and its plans for the near future.

As reported in our May and September issues, Version 2 of SNMP has been “in limbo” awaiting the resolution of issues relating to security. An advisory team was formed to analyze the two existing SNMP security specifications (USEC and SNMPv2*). This month we bring you the team’s first status report.

And with that we come to the end of 1996 and the end of ten years of *ConneXions* in printed form. I hope you will join us as we re-emerge on the Web and that you will help us create an online journal which continues to track current emerging standards and technologies in the computer and communications industry. See you online!

Internet Digital Libraries

by Jack Kessler, kessler@well.sf.ca.us

Introduction

Anyone—certainly any printed journal—in possession of an archive and an information channel as rich as those of *ConneXions*, will be looking nowadays at both print and digital media. Text and images and many other things besides all can be stored and searched and retrieved and used now, in print and online and in perhaps too many different ways.

Digital Libraries and the Babel of the Internet

The *Digital Libraries* term is in use today to describe projects which come in a variety of flavors: “about 400,000” entries for it in Internet indexes like AltaVista [1] justify some curiosity as to just what the Digital Libraries term actually means:

- *Systems projects*: There are vast government-sponsored systems design efforts, such as the six projects of the US National Science Foundation’s *Digital Library Initiative* [2], or the EU’s *Tele-matics for Libraries* and *INFO 2000* programs [3];
- *Computer projects*: There are very many ongoing efforts under way within the computer science community to improve the organization of digital information used by databases and computers generally. Oracle [4] and Informix [5] and Sybase [6] are three leading among the many commercial firms trying continually to define new storage and retrieval methods for digital objects;
- *Information projects*: Digital Libraries are said to exist nowadays wherever and however digital information is being assembled. Projects range from detailed academic studies compiling meticulous statistics and images and other data, to print publishing industry ventures, to the vast assortment of institutional and personal “e-conference archives” and “Web sites” which, each pursuing its own often eccentric character and structure, all now choke the Internet’s bandwidth;
- *Library projects*: There are, finally, many efforts being made by the traditional “library” information community, worldwide, to come to grips with—to marry their older professional skills with—the new techniques of digital information. Some of these “library” Digital Libraries projects are online, some are not; some look like the other, computer-and-systems-based projects, others look simply like printed books and catalog cards which suddenly and simply have “appeared” on computer screens.

These four categories themselves might be multiplied endlessly. There are that many Digital Libraries projects—both so-called and as named by others—which are operating already and which rapidly are growing in number. There are Digital Libraries in Thailand and Australia and Japan, and Digital Libraries projects under way ranging from local efforts to organize slide collections to international work on the collections of the Vatican Library. [7]

The effort to understand what all of this has in common is the effort which interests me personally the most.

There is an old saying in academia that “If something is everything then maybe it’s nothing.” This has become increasingly applicable to many aspects of the digital information revolution.

The range of interests in digital information is in fact enormous. The World-Wide Web is not the whole Internet, and the Internet is not even William Gibson's entire *Matrix*; HTML is not the whole world of marked-up text, and SGML is not the entire picture in the digital representation of information. There are so very many projects, so many of them unaddressed by current standards efforts and even by current terminology.

There is a growing, crying, need to discriminate and to define. There seems to be a real need for defining something—some element in common, hopefully—at which all of the various Digital Libraries efforts are trying to grasp. The term never would be in use—in common, in so many different places—if a lot of people did not at least sense that there was some common good purpose which might be served.

Data versus Information

To define this general Digital Libraries purpose better, another distinction is useful, that between *data* and *information*: information as data organized and presented so that users can use it—information as the bits with knowledge added, data as just the bits without the brains.

The outstanding current digital information problem arises from the fact that today's Internet, to today's newest digital information user—the *general public*—is just data: just the raw material without enough "value added," just the bits without the brains. If information is defined as bits which have been organized and presented "so that users can use it," then today's information-overloaded, on-again/off-again, hypertext loop-plagued Internet, is not "information": at least it is not very usable to the new general public users, who never have been online or even near a computer before.

New markets: Users, Clients and Customers

This was not so much of a problem back when all Internet users were engineers possessing vast computer knowledge and experience. But since the "acceptable use restriction" gloves were taken off, in 1992, a whole new general public market—that 63% of US households which still do not have a computer, and even a sizable majority of the households which do—has begun to discover the delights and confusions of online digital information. Organizing and presenting for this general public market is a whole new exercise, very different from the same effort formerly aimed primarily at computer engineers.

Digital Libraries currently, then—the 1) Systems projects, the 2) Computer projects, the 3) Information projects, and the 4) Library projects—all might best be viewed in light of their responses to this latest challenge, that of converting online digital data into information for this new general public group of users.

This is not so much the older problem, any longer, of converting other types of information into data; nor is it the problem of the "storage and searching and retrieval" of that data. It is the problem of getting the data—eventually, somehow—into information formats acceptable to the entirely new class of general public users.

The "value added" needed for general public users is less technical: "images" and "links" and "sound" of course, rather than just dumb printed text installed online—but it has less to do generally now with the technology, on which so much successful work already has been done, and more to do with the users themselves, with the approaches and psychology of sales and marketing and customer service and professional assistance.

Internet Digital Libraries (*continued*)

One principal architect of a leading Digital Libraries project refers to the “glue” which holds digital information together (Stanford’s Terry Winograd) [8]. This glue no longer is one of the high-grade, overly-sensitive, and very-expensive adhesives of the Internets earlier test-bed era. Now, in the coming decade of America Online and Network Computers and Netscape and @Home and TCI and Viacom, we are talking Elmer’s—lots of it.

The shift has been from research and applications development which used to focus upon the technology, to research and applications development which now focuses upon the users. Technology enthusiasts should be happy. Xerox PARC [9] has been preaching for some time that digital technology would become successful only as it succeeded in becoming 1) ubiquitous, 2) inexpensive, and 3) invisible: i.e., found everywhere, assumed to be useful, taken for granted, like the telephone and the television and the toaster. The day has arrived, for the general public, at least.

It makes basic marketing sense. You can sell more units and services to more people this way. You can lower margins and raise volume, and realize the marketer’s magic *economies of scale*.

Something new: The uninterested General Public User

The key to any marketing, though, is to understand the user—the client, the customer—thoroughly. This was not so much of a problem back when digital information was used only by engineering students and professionals. Back then there was a single user profile, and one which was fundamentally friendly toward the technology. Most engineers, faced with a computer and an information system, were fascinated, and wanted very much to learn to use it all no matter what it did.

This is not the case now, though, with the new general public users. They only want the “information.” Yesterday’s Internet users were in love with the idea of the Internet, nearly regardless of what information—or data—which it might or might not contain. Today, that interest no longer is there. Today’s users do not “want to know how the car works,” they just “want to drive it.”

It is not that general public users are less intelligent, or even less-educated, than their computer engineering forebears were. It is just that they have other interests: car repair, stamp collecting, changing a diaper, going to the beach—they are busy with those. Since the demise of acceptable use restrictions, online digital information increasingly is having to address an entirely new phenomenon: the *uninterested* user.

Today’s Digital Libraries—all of them, the 1) Systems projects, the 2) Computer projects, the 3) Information projects, and the 4) Library projects—are designing for this un-interested general public user, or at least the more up-to-date ones among them are.

“R” going one way, “D” going another?

This has distressed some members of the online digital information community. It may even have caused a rift. There are plenty of sophisticated applications under development—vast numbers of high-bandwidth and otherwise-expensive ideas and projects—which would be entirely derailed by a total migration of online digital information to a world populated only by America Online and Network Computers and Netscape and @Home and TCI and Viacom.

As UC Berkeley's Stephen Cohen says, "People forget that in 'R&D', companies don't do the 'R', they only do the 'D'" [10]: people also forget that we would not have the "D"—that present world of America Online and Network Computers and Netscape and @Home and TCI and Viacom—if it had not been for the Internet testbed "R" which preceded it.

Proposals under way for a high-speed, research-oriented Internet II [13]—already being called "Son of Internet" by some, "Grendel" by others—indicate that what may emerge may be an online digital information split, with high-speed research applications going one way, and the less-expensive and less-capable general public market going the other.

But Digital Libraries can help in the higher-end efforts as well. The MBone, one of the more promising high-bandwidth transmission projects, already has its catalog—of past and future transmissions—under development online, and its archive, and all of the attendant problems of the categorization and classification and indexing and abstracting and search and retrieval of same by users [11].

Similar questions come up with the Internet's URNs and URLs and domain names, and with the proliferating SGML DTDs, and with the W3 indexing META hidden tag system which looks increasingly like the old MACHine Readable Cataloging/MARC format used for years for printed books.

These—catalogues and archives and categorization and classification and indexing and abstracting and search and retrieval—all are traditional library questions. They were questions asked in the past about illuminated manuscripts and about printed books, and they are being asked now about digital "documents" and online information. They have less to do with the "digital" side of the Digital Libraries equation, than they have with the fuzzier, less clearly-defined, "libraries" side.

They have to do with users, and with what it takes to make data usable to a user as information: whether this is data recorded in ink on parchment or registered as bytes in a bitstream, and whether the eventual information is to be used by "high-tech researchers" or by members of the general public.

Digital Libraries, then, is somewhat of a misnomer. As used currently, it describes too many things. But the term describes so many things that it must describe at least some one thing which they all have in common. That something, I suggest, is the conversion of data into information, the latter being "whatever is useful to the current group of users."

One great challenge of the 1990s is that this "current group of users" suddenly has exploded out to include not only the traditional specialists but a general public which is un-interested in the underlying technique. Digital Libraries methods for coping with this challenge can help meet the demands of more technically sophisticated and/or interested users as well.

**Something else new:
The international
uninterested
General Public User**

One other great late-1990s challenge for digital information, then, is that users very suddenly are located in many nations around the world.

Internet Digital Libraries (*continued*)

There are radio modems in use suddenly in Cambodia; China is online; Mozambique just acquired connectivity; the Internet's 9+ million host count jumped to 12+ million just in the first six months of 1996 [12], and great percentage growth figures now may be found outside the US. I send out an e-journal myself, every month, to readers in over 70 countries, which reaches all of them in milliseconds: a publication and distribution miracle which must have W.H. Smith and W.R. Hearst and H.R. Luce turning in their graves.

The great challenge of the coming decade is going to be that of ensuring not that all sorts of users everywhere will have access to digital information—this problem rapidly is being solved now—but that the inundation of such information will not be such that users split into “high-end” and “low-end,” “digital knowledge” and “digital ignorance,” “the digitally empowered” and “the digitally disenfranchised.”

It all will be digital: “print” already is—“photography” and “sound” and “TV” and “cinema” and “telephony” all are getting there. There are many remaining technical challenges and problems, from distributed processing and the scalability of high speed transmission, to multilingual techniques and the development of object relational databases. But increasingly now the key problem is not how to digitize, but how to organize and present whatever is digitized to users. This is the fundamental problem-in-common to which all Digital Libraries are dedicated.

Internet Digital Libraries: The international dimension

There is a lot of Digital Libraries work under way now, in a lot of places, all of it filling at least the four categories of 1) Systems projects, the 2) Computer projects, the 3) Information projects, and the 4) Library projects suggested here. I just now am publishing a book which tries out a beginning overall view:

Jack Kessler, *Internet Digital Libraries: The International Dimension*, Artech House, 1996, ISBN 0-89006-875-5. [14]

The book does not address the Digital Libraries problem so much theoretically or philosophically, as I have done a bit here, as it does internationally: it is filled with examples of current Digital Libraries work under way in places like Chiang Mai, Thailand, and Surabaya, Indonesia, and Lyon, France—coupled with the suggestion, implied throughout and declared directly whenever I can, that the development of any Digital Libraries solutions now will be/will have to be internationalist in its approach.

A great deal has been written on the general subject, however. There are some such references given in my book. Other materials by me and by others, and references and live links and nice pictures and even bibliographies and resource lists, may be found at:

<http://www.fyifrance.com>

The general problem of Digital Libraries has been dealt with recently by thinkers and writers as diverse as Blaise Cronin, Michael Buckland, Walt Crawford, Clifford Lynch, Kenneth Dowlin, Michel Melot, Jesse Shera and Wilfred Lancaster.

If you embrace the broad definition which I am encouraging—that the conversion of “data” into “information” is what is involved here—there are resources of interest to you about the 15th century transition from whatever-preceded-it to print (Elizabeth Eisenstein), and about any time and place where one mode of expression has been succeeded by or even simply influenced by another (Walter Ong, Jack Goody, Pierre Levy, Marshall McLuhan, Henri-Jean Martin, Erich Auerbach, etc., etc...).

There is no real need to reinvent the wheel on all of this. It has not *all* been invented already in the past and elsewhere, but there is a lot which can be learned—and money which can be saved—by looking. Digital technology may be new, but its newest un-interested general public user has not changed much in a very long time. The Digital Libraries problem no longer is whether to choose the new digital medium over traditional print, but to choose it, and use it, effectively for the users.

References

- [1] <http://www.altavista.digital.com>
- [2] <http://www.grainger.uiuc.edu/dli/national.htm>
- [3] <http://www2.echo.lu>
- [4] <http://www.oracle.com>
- [5] <http://www.informix.com>
- [6] <http://www.sybase.com>
- [7] <http://witloof.sjsu.edu/courses/281.kessler/sjsu210.htm>
- [8] <http://www-pcd.stanford.edu/hci/people/winograd.html>
- [9] <http://www.parc.xerox.com>
- [10] <http://sunsite.berkeley.edu/IS3/fyi3.html>
- [11] <http://www.msri.org/mbone/>
- [12] <http://www.nw.com>
- [13] The “Internet 2” articles in this issue of *ConneXions*.
- [14] http://www.artech-house.com/artech/html/catalog/book_html/0-89006-875-5.html?d8j9h8
- [15] Barron, Billy, “Another use of the Internet: Libraries Online Catalogs,” *ConneXions*, Volume 5, No. 7, July 1991.
- [16] Lynch, Clifford, “Linking Library Automation Systems in the Internet: Functional Requirements, Planning, and Policy Issues,” *Library Hi Tech*, Volume 7, No. 4, 1989.

JACK KESSLER, has been an Internet Trainer and Consultant since 1991. He studied philosophy, politics, law, and library and information science at Yale, Oxford, and the University of California, and spent 15 years in the importing business in East and South Asia, the US and Europe. Since 1992 he has edited *FYI France*, ISSN 1071-5916, a monthly e-journal of digital information news about France and Europe. He writes frequently on subjects ranging from the Minitel and the Internet and W3 to libraries and librarianship, printing history, and the information professions. Beginning in 1997 he will edit the *FYI France^{SM TM} Online Service*, <http://www.fyifrance.com> (already online), an experiment in combining W3/passive and e-mail/active marketing techniques in the commercial provision of on-line information service. He is nearly impossible to reach except via e-mail at: kessler@well.sf.ca.us

The Focus of the Internet-II Vision

by David L. Wasley, University of California

Introduction

The parable about the Blind Men and the Elephant applies very well to the current variety of views on what the term "Internet-II" means. Some in the higher education community see it as a separate network exclusively for higher education; for others it is viewed as a research testbed; still others believe it is a publicity ploy to get more Federal money for higher education; many outside of higher education think it must be a solution for the problems of the current Internet. In the view of many of us involved in the planning, there is some degree of validity in each of these views but they all miss the essential point of the initiative. This article will try to convey the important fundamental nature of the Internet-II initiative as seen by the information technology community that has developed and nurtured it.

It is undeniable that the quality of service as perceived by users of the existing commodity Internet has deteriorated over the last several years. This has had a serious impact on researchers and scholars who must rely on the Internet for their work, must transfer large amounts of data between remote sites, or who try to use streaming applications such as packet video for conferencing or collaboration. However, for the majority of users the occasional delay in delivery of a Web screen is more of a frustration than a strong deterrent. Thus, although the higher education community is concerned with this problem, it alone would not have generated the strong consensus and momentum behind developing a greatly enhanced Internet infrastructure.

Driving forces

The higher education community has incorporated information technology into its academic and research programs for over 30 years. It is not uncommon today for a campus, particularly in the research community, to have a high speed ubiquitous network and at least one computer on most faculty, staff and graduate student's desks. The Internet itself grew up in this environment long before it became popular in the commercial world.

This long familiarity with the use of information technology continues to inform and expand the vision of how such technology might transform teaching and learning, scholarly collaboration and research, and ultimately the university itself. Today teachers are able to have far greater contact with students through the convenience of e-mail. Class assignments and background materials can be found on the Web so there is no longer the worry about misplacing those essential pieces of paper. Collaboration between students and among researchers takes place over the network asynchronously and far more effectively than might have been achieved before. A number of experiments have shown that it is possible to deliver an entire course in certain fields remotely over the Internet. Yet all these enhancements are merely scratching the surface of a world of compelling possibilities.

The IT leaders of the major US research universities, in response to increasing demands from teaching and research faculty, have articulated a common vision of what is needed in order to realize the next generation of network-based applications of importance to the higher education community. They include new modes of interactive collaboration and new ways of enhancing learning through information technology, the integration of complex distributed digital library collections with academic programs, greater access to specialized research facilities such as accelerators and supercomputers, and life long scholarly pursuits facilitated through ready access to learning materials from homes, offices, or anywhere convenient to the learner.

Electronic access to a broad range of new types of information as well as “raw data” will engender new possibilities for the synthesis of ideas that can ultimately lead to the creation of important new knowledge.

The university as a center of knowledge and learning must become integrated more fully into the fabric of industry and commerce. Modes of learning and collaboration developed during periods of residential education, our traditional period of introduction to higher education, can be perpetuated if sufficient access to learning resources is available. A basic component of this will be an ubiquitous advanced communications infrastructure that can enable teaching and learning to remain a part of everyday life. The future American work force thus equipped will continue to vitalize industry and commerce and help to maintain our world leadership in critical emerging fields.

Why the emphasis on an advanced network?

It is this vision of new modes of higher education enabled through a set of advanced distributed information systems that motivates and energizes the higher education information technology community. The realization of this vision will require major advances in pedagogy, supporting resources, computer software and technology infrastructure. A fundamental component of this complex picture, without which many of the other pieces can not even be developed, is the availability of a rich set of communications services in support of advanced applications requirements.

In the commodity Internet today, there is exactly one “quality of service”—best effort delivery of most data packets at some unpredictable time to a single destination. All current applications are designed with this limitation in mind. The new advanced applications that are envisioned will require greater guarantees of predictable delivery, within a defined time window, and with an acceptable transport capacity. Some applications will require the delivery of the same data to many recipients. Some will tolerate asymmetrical service quality and others will require closely symmetrical service. Each application must be able to request and acquire the required services (or wait until they are available) dynamically.

It is not at all certain that the commercial commodity Internet is prepared to move quickly to deploy the range of advanced Internet communication services that higher education seeks. Higher education now represents only a small fraction of the Internet consumer market. A successful business must focus on developments that “maximize the bottom line” and for the Internet market this means the broad consumer customer base and the high end commercial and industrial customers. These customers are not (yet) demanding the types of data communications services that are needed to support the new types of applications envisioned by higher education.

Given this perspective, information technology leaders in the higher education community have come together with partners from the Federal government and industry to develop a joint strategy to catalyze and accelerate the development of “next generation” applications and the infrastructure required to enable them. The focus of this broad effort is on realizing the vision described above. The first, and in many ways easiest, step is to cause to be put in place a new set of advanced Internet communications services and make them readily available to the community of applications developers and users.

The Focus of the Internet-II Vision (*continued*)

This advanced communications infrastructure has come to be called "Internet-II" because it builds on the common bearer service of the existing Internet—the Internet Protocol—with extensions and enhancements now in the later stages of development in labs and vendor shops. Unfortunately the same terminology has been applied to the broader goals of the applications space as well which obscures the primary meaning of the Internet-II initiative.

The Internet and Internet-II

Concern has been expressed that higher education is abandoning the commercial Internet industry. Far from it. It is estimated that more than 70% of the Internet traffic entering or leaving our campuses is exchanged with non-university sites. The advanced Internet infrastructure will serve only as a proving ground among those institutions developing the new applications. Furthermore, it is anticipated that success within this pre-competitive arena will lead to early adoption of the new technologies and service delivery models by commercial Internet providers. Only the private sector that can build and support the infrastructure required to bring the advanced services into homes and offices to reach the ultimate beneficiaries of the new applications. Therefore it is clear that higher education not only will continue to rely heavily on the commercial Internet for most of its communications services today but is planning for technology transfer to the private sector as soon as possible.

The enhanced Internet architecture, as currently conceived, includes a new infrastructure component that will enable the higher education community to continue to make use of existing Internet services as well as gain access to the new advanced Internet-II services. This component has been termed the "GigaPOP" but might be described more usefully as an *Internet-II Services Center* (I2SC). One might think of an I2SC as the Internet equivalent of a regional shopping mall: a vendor-neutral place where consumers in the region can come to acquire a variety of Internet transport services. A campus or research institution will be able to lease a single high capacity connection to an Internet Services Center at which it can gain access to differentiated services from one or more pre-competitive as well as commercial service providers. This should lead to healthy competition among Internet service providers as well as prove to be a highly cost effective way for Internet consumers to acquire services.

It must be emphasized that although the I2SCs will be able to make available a wide variety of Internet services, the Internet-II Project community will be defined as those institutions that not only connect to the I2SCs but also provide campus infrastructure support for end-to-end advanced Internet services between participating research and user locations. However, a very important aspect of proving the value of the new applications will be to observe their efficacy in a variety of school, community and professional environments. Therefore it is expected that the community of participants will expand as the project matures and applications become available for testing among a wider audience.

The Project

The Internet-II Project is focused on developing and proving the effectiveness of a wide range of new teaching, learning, research and collaboration technologies. Participants will commit considerable financial and intellectual resources towards the successful outcome of this project. The project also will require the availability of advanced Internet services and much higher capacity and "intelligence" in the network than exists today. Achieving this infrastructure will require continued close cooperation of the entire Internet-II Project community.

Seed money will help move this vision towards reality but ultimately the result must be self-sustaining and commercially viable. Higher education looks forward to working closely with partners in the information technology and communications service industry to achieve these goals in the most cost effective and broadly applicable manner.

DAVID WASLEY holds a Masters Degree from the University of California, Berkeley, and has been a member of the staff of the University for 27 years. For the last decade he has been responsible for the development of the UC Berkeley campus data network and associated services. He was active in the founding and development of the Bay Area Regional Research Network (BARRNet). He is co-author of RFC 1709 "K-12 Internetworking Guidelines." Recently he joined the UC Office of the President in order to focus on new issues and challenges in the area of Information Infrastructure Planning. E-mail: David.Wasley@UCOP.EDU

Internet-II Architecture

Introduction

This technical overview of the proposed architecture for the Internet-II project was created by a working group composed of Scott Bradner (Harvard Univ.), Scott Brim (Cornell Univ.), Steve Corbato (Univ. of Washington), Russ Hobby (Univ. of California, Davis), and David Wasley (Univ. of California System), with contributions from many other individuals, including in particular a presentation by Professor Larry Landweber of the University of Wisconsin at a workshop in Ann Arbor, Michigan, in July 1996. It is intended to be a working document and will be updated as the Internet-II project evolves.

Internet-II Project

The Internet-II project is a collaborative effort among a number of universities, federal R&D agencies, and private sector firms to develop a next generation Internet for research and education, including both enhanced network services as well as the multimedia applications which will be enabled by those services. The work is developmental and pre-competitive in nature. It is more fully described in a companion document to this technical overview (see page 14).

Objectives

The technical objectives of Internet-II include:

- Maintain a common bearer service to support new and existing applications
- Move from best effort packet delivery to a differentiated communications service
- Provide the capability of tailoring network service characteristics to meet specific applications requirements
- Achieve an advanced communications infrastructure for the Research and Education community

Applications requirements

In a number of technical meetings and workshops over the past several years, faculty members and other university representatives have identified a set of advanced applications that will greatly enrich teaching, learning, collaboration and research activities.

A major impediment to the realization of these applications is lack of advanced communications services. The broad use of distance learning will require selectable quality of service and efficient "one-to-many" data transport in support of multimedia and shared information processing. To support world class research on a continuing basis, the academic community requires high capacity and selectable quality of service to make effective use of national laboratories, computational facilities and large data repositories.

continued on next page

Internet-II Architecture (*continued*)

Network services

Internet-II is designed to provide a variety of services “on demand” in support of advanced applications. These dynamically selectable services will include guaranteed bounded delay, low data loss, and high capacity. For example, in order to support delivery of advanced multimedia teaching materials from a digital library repository to a dispersed audience of learners, it will be necessary for the service delivery infrastructure to support “multicast” data delivery with guaranteed upper bounds within the transport components on delay and data loss.

New protocols to enable this functionality have already been defined and will be deployed early in the Internet-II project. These protocols include the IETF defined quality of service protocols such as RSVP and RTP along with IPv6, the IETF-developed replacement for the version of IP that is in current use on the Internet. In addition, Internet-II will provide access to the underlying network infrastructure for those environments that can support that access and for those applications that can make use of specific capabilities offered by the infrastructure.

Implementation

At the heart of the Internet-II design is a new technology for providing advanced communications services. The technology, referred to as a *GigaPOP*, is a complex of technologies developed over the first decade of the Internet integrated with new technologies developed by vendors and the Internet Engineering community. The Internet-II project will demonstrate proof of concept of this new set of technologies and services so that they can become part of the next generation of commercial Internet service offerings.

A GigaPOP is the point of interconnection and service delivery between one or more institutional members of the Internet-II development project and one or more service providers. Typical institutional connections will be made via ATM or SONET services at very high bandwidth. The fundamental advance represented by the GigaPOP architecture is dynamically acquired “quality of service” in support of a broad range of new applications while maintaining a common interoperable “bearer service.” Service characteristics will include end-user definable capacity as well as latency. An essential part of the Internet-II project will be to determine the incremental costs associated with support of differentiated classes of service and to develop the mechanisms to collect data about the use of these resources by individual users.

The architecture of the GigaPOP also will support service delivery to regional or state-based not-for-profit consortia such as the Virginia Educational Network, the Washington State K-20 network, or the combined University of California and California State University system. It is envisioned that 20-30 GigaPOPs nationwide will comprise the Phase 1 deployment. These will be designed and managed collectively on behalf of the Internet-II project community.

Equipment at a GigaPOP site will include:

- One or more very high capacity advanced function packet data switch/routers capable of supporting at least OC-12 (622 megabit/second) link speeds and switched data streams as well as packet data routing;
- Switch/routers supporting Internet Protocols (both version 4 and the new version 6), advanced routing protocols such as MOSPF, and “quality of service” protocols such as RSVP;

- SONET or ATM multiplexers to enable allocation of link capacity to different services such as highly reliable IP packet delivery, experimental testbeds for emerging protocols, or special requirements determined by new initiatives among the Internet-II member institutions;
- Traffic measurement and related data gathering to enable project staff to define flow characteristics as part of the operational and performance monitoring of the GigaPOPs.

One or more wide area communications service providers will connect to the GigaPOPs in order to provide communications paths between the nationwide set of GigaPOPs and between GigaPOPs and the established commercial Internet. Thus, participating institutions will be able to acquire a wide variety of commercial as well as pre-competitive communications services over a single high capacity communications link to the nearest GigaPOP facility. In particular, to support high performance distance learning and remote collaboration initiatives, the GigaPOP architecture will facilitate local interconnectivity between the higher education community and those commercial providers offering emerging high-bandwidth home access technologies.

The most advanced applications will require a set of communications paths among the GigaPOPs that are engineered especially for the Internet-II project. It is essential that these interconnect pathways fully support the protocols and functions noted above. Recently, NSF has announced a *High Performance Connections* grant program which expands its vBNS infrastructure to connect as many as 100 sites nationally to the current OC-3/OC-12 backbone and could provide a deployment platform for emerging applications in support of research and collaboration. It is envisioned that the HPC/vBNS, with its new capabilities, will be the initial interconnect network among the GigaPOPs. If the vBNS should prove insufficient for the full range of Internet-II requirements, other alternatives will be employed.

Although direct SONET pathways might be most effective in providing the inter-GigaPOP pathways, it seems most likely that ATM-over-SONET will be the most commonly available commercial service. Because Internet-II will use virtual connections within and between the GigaPOPs, a test network can be implemented alongside of the production network without having to duplicate facilities.

This test network will be used to experiment with new capabilities of the network itself where the production network can be used to provide reliable service for applications.

Operations

Clearly, the design of the GigaPOPs must meet the requirements of very high reliability and availability. Each GigaPOP site will be physically secure and environmentally conditioned, including backup power and resistance to damage from acts of nature. Physically diverse fiber optic and wireless communications paths will maximize service robustness against the unlikely event of physical damage external to the site. In addition, the Internet-II infrastructure will be designed to be secure from the threats of those who would seek to disrupt its operations.

Not all GigaPOP sites will be staffed 24 hours per day. Instead, redundant Network Operations Centers will monitor the operation of all equipment remotely via both in-band and out-of-band circuits and will dispatch problem resolution staff as needed to effect restoration of normal services.

continued on next page

Internet-II Architecture (*continued*)

Conclusion

The Internet-II architecture has been chosen to demonstrate the effectiveness of new technologies in providing the next generation communications infrastructure. The success of Internet-II will allow higher education and research institutions to remain world leaders in the development of advanced applications of information technology.

Scope of the Internet-II Project

Introduction

Building on the tremendous success of the last ten years in generalizing and adapting research Internet technology to academic needs, the university community has joined together with government and industry partners to accelerate the next stage of Internet development in academia. The Internet-II project, as it is known, will bring focus, energy and resources to the development of a new family of advanced applications to meet emerging academic requirements in research, teaching and learning.

The project will address major challenges of the next generation of university networks. First and most importantly, a leading edge network capability for the national research community will be created and sustained. For a number of years beginning in 1987, the network services of NSFNET were unequaled anywhere else. But the privatization of that network and the frequent congestion of its commercial replacement have deprived many faculty of the network capability needed to support world class research. This unintended result has had a significant negative impact on the university research community.

Second, network development efforts will be directed to enabling a new generation of applications that fully exploit the capabilities of broadband networks—media integration, interactivity, real time collaboration—to name a few. This work is essential if new priorities within higher education for support of national research objectives, distance education, lifelong learning, and related efforts are to be fulfilled.

Third, the work of the Internet-II project will be integrated with ongoing efforts to improve production Internet services for all members of the academic community. A major goal of the project is to rapidly transfer new network services and applications to all levels of educational use and to the broader Internet community, both nationally and internationally.

The project will be conducted in phases over the next three to five years, with initial participation expected from one hundred universities, a number of federal agencies, and many of the leading computer and telecommunications companies. A summary of the technical architecture is contained in a companion document, entitled “Internet-II Architecture” (see above, page 11).

In the initial project phase, end-to-end broadband network services will be established among the participating universities. On a parallel basis, applications design will commence using teams of university faculty, researchers, technical staff and industry experts. It is expected that within eighteen months, “beta” versions of a number of applications will be in operation among the Internet II participating universities.

University participation in Internet-II

At a meeting in Chicago in October, 1996, representatives of thirty-four universities agreed unanimously to endorse the goals of the project, committed their institutions to finding the resources necessary to participate in the project, and pledged initial funding to enable planning efforts to proceed without delay. Since that time, the number of participating campuses has reached seventy, and further additions are anticipated.

Internet-II charter membership is open to institutions who join the project prior to the general membership meeting in January, 1997.

Joining the project involves the following basic institutional commitments:

- Creating a project team to support the applications development and advanced network services objectives of the project.
- Establishing broadband Internet connectivity on an end-to-end basis as soon as possible to support development, testing and use of applications
- Participating at the executive level in the overall management of the project.
- Contributing necessary financial support to the above activities and to the central management and administrative expenses of the project.

Based on the information available at this time, it is estimated that institutional expenses of all kinds may reach \$500,000 per year over the next several years. A substantial fraction of this amount may be covered by networking and related budgets already in place, depending on individual institutional circumstances. An additional commitment of up to \$25,000 per year is required to defray the central administrative and member support expenses of the project.

Internet-II partnership and funding arrangements

As announced by President Clinton on October 10, 1996, the federal government will participate in Internet-II through the programs of its major research agencies, who will provide grant support in their areas of program interest, such as the NSF High Performance Connections initiative. In most respects, the partnership and funding arrangements for Internet-II will parallel those of previous joint networking efforts, of which the NSFNET project is a very successful example. Industry partners will work with campus-based and regional university teams to create the advanced network services that are necessary to meet the requirements of broadband, networked applications.

Funding for Internet-II will include both financial and in kind services and products of various types that will be necessary for the project. Since most of the project effort will occur on or near university campuses, it is anticipated that the majority of funding from government research agencies and industry partners will be in the form of grants to the participating universities.

References

- [1] NNSC Staff, "Profile: NSFNET," *ConneXions*, Volume 1, No. 2, June 1987.
- [2] Hans-Werner Braun, "The new NSFNET backbone network," *ConneXions*, Volume 2, No. 12, December 1988.
- [3] P. Bosco & H.-W. Braun, "The NSFNET T1/T3 Network Celebrates 5 Years," *ConneXions*, Volume 7, No. 9, September 1993.
- [4] Harris, S. & Gerich, E., "Retiring the NSFNET Backbone Service," *ConneXions*, Volume 10, No. 4, April 1996.

A comparison of the performance of 100-BaseTX and 100VG networks

by Roger Cohen, Netmarq

Introduction

100VG and *100-BaseT* are two different local area network (LAN) technologies that operate at a rate of 100 Megabits per second (Mbps). (See “100VG and 100-BaseT Tutorial” in the November 1996 issue of *ConneXions*.) They are competing to become the de facto standard for new LANs and for upgrades to existing Ethernet LANs, and represent two contrasting solutions to the problem of upgrading the original 10 Mbps Ethernet standard to operate at 10 times that speed. Both technologies provide a shared medium for interconnecting workstations and servers at 100 Megabits per second over unshielded twisted pair cable (UTP). They both use exactly the same packet format as 10-BaseT Ethernet, so they can be transparently bridged to 10-BaseT Ethernet. They differ in their media access method and in the standard of UTP cable required.

Key findings

These two technologies were compared using the *Netmarq Network Performance Test Suite* in a NetWare 4 environment with up to 18 workstations and 3 servers. The test suite measures real data throughput between network clients and servers. Four makes each of 100-BaseTX network interface cards and repeaters were tested in various combinations and compared with Hewlett-Packard 100VG cards and repeaters. Key findings were:

- 100VG and 100-BaseTX could both operate at the maximum expected rate when the data flow was exclusively from server to clients.
- Only 100VG operated at the theoretical maximum rate with equal data flow in both directions.
- Two data patterns designed to cause “baseline wander” in 100-BaseTX resulted in large reductions in 100-BaseTX throughput, the sizes of which were sensitive to the lengths of the leads between the repeater and the clients and servers.
- 100VG generally showed less variation in successive repeats of the same data transfer activity than 100-BaseTX.

The technologies

100-BaseT is designed to be as similar as possible to 10-BaseT Ethernet and handles contention between multiple devices trying to transmit simultaneously using exactly the same collision detection method (CSMA/CD). A device with data to transmit listens to the shared wire until it is quiet then transmits its next data packet. If another device attempts to transmit data at the same time a *collision* is detected by both devices and they each wait for a pseudo-random period before attempting to retransmit. The back-off period needs to have an element of randomness to ensure that each device waits for a different period before retransmission to avoid a chain of successive collisions.

100-BaseT's use of the existing Ethernet collision detection mechanism reduces the maximum allowable size of a single collision domain by a factor of 10 compared with 10-BaseT to a diameter of about 250 metres on copper cable. As with 10-BaseT, 100-BaseT can be run on several different physical media. The tests reported here were performed on the predominant 100-BaseTX form which uses two pairs of a category 5 UTP cable (see below for a more detailed description of 100-BaseTX).

In 100VG a *Demand Protocol* access method is used. A device signals to the 100VG hub when it is ready to transmit data while the hub continuously polls only those devices with data to transmit, allowing each one in turn to transmit a certain amount of data. 100VG uses all four pairs in a standard UTP cable and will operate over either category 3 or category 5 cable.

The main practical differences between the technologies in terms of installation lie in the standard of cabling plant required. 100VG can usually operate as a direct replacement for 10-BaseT Ethernet using exactly the same cable and replacing 10-BaseT repeaters with 100VG hubs, with the proviso that all four pairs of the UTP cable are available. 100-BaseTX has much more stringent requirements. All cable runs must be fully category 5 compliant from end to end, and a maximum of two repeaters is allowed in a single collision domain.

Baseline Wander

Conventional 10 Megabit per second Ethernet has several well-defined physical implementations involving different types of cable; thick coaxial cable (10Base5), thin coaxial cable (10Base2), unshielded twisted pair (10BaseT), or fiber optic cable (FOIRL). Similarly, 100-BaseT has several different physical manifestations. 100-BaseTX uses 2 pairs of category 5 UTP; 100-BaseT4 uses 4 pairs of category 5 UTP; 100-BaseTF uses fiber optic cable. The overwhelmingly predominant implementation is currently 100-BaseTX, and that is the one we tested.

There is however a theoretical deficiency in most implementations of 100-BaseTX. When data is to be transmitted it is first converted into a series of binary digits (0s and 1s). The MLT-3 encoding scheme used in 100-BaseTX converts these binary digits physically into one of three voltage levels applied to the wires; -1, 0, and +1 volts. For every binary 1 transmitted, the voltage changes to the next level in the sequence -1, 0, +1, 0, -1... For every 0 transmitted, the voltage stays at its existing level. Certain data patterns cause the 100-BaseTX transmitter to send out a long sequence of successive zeroes which means that it is possible for the signal to assume a constant level of +1 volts and remain there for a relatively long period. The 100-BaseTX transmitter contains a component (a transformer) which is intended to protect it from being damaged by large voltages on the cable, and in some designs this saturates with current and fails to keep the voltage at the +1 level for the required period, instead allowing it to drop towards 0 volts. When the sequence of zeroes ends and more voltage transitions finally occur, the signal swings from 0 to -2 volts instead of from +1 to -1 volts. Gradually, it moves back to the +1 to -1 volt swing as the current in the transformer dies away. This is called *baseline wander* and it causes loss of data at the device that is receiving the signal.

Hewlett-Packard engineers have devised a publicly-available set of data patterns that are expected to cause differing degrees of baseline wander. The first (Pattern 1) is 40 bytes long and causes sub-maximal baseline wander. The second (Pattern 2) is 990 bytes long and is intended to cause maximum possible baseline wander. We adapted our standard performance test to use each of these patterns, as explained below.

The performance tests

Netmarq was commissioned by Hewlett-Packard to compare 100-BaseT with 100VG in terms of speed, stability, ability to run successfully on maximum length cables, and susceptibility to data patterns that cause baseline wander.

Performance of 100-BaseTX and 100VG (*continued*)

We constructed several test networks with two or three NetWare 4 file servers, 22 workstations, and one or two 100-BaseT hubs. The cables from the workstations and servers to the hubs were designed to emulate a standard category 5 UTP installation. Each cable run was formed from three sections: two short patch cables connected to each end of a 90-metre run of solid cable. The total length of each run was 97 metres and every run was certified as a whole by a cabling contractor to be category 5 compliant.

In order to measure pure network performance we normally use two tests from our standard test suite in which the workstations read or write each 30 Kilobyte record alternately from a database file containing only two records. Both tests put very heavy loads on the network. In the *Overlaid Read* test each workstation reads alternate 30 Kilobyte records from a shared file containing only two records. The file is fully cached in the file server, so this test measures throughput from server to workstation without involving the server's disk subsystem. In the *Overlaid Read/Write* test each workstation reads and writes alternately to two 30K records in a private file—not a shared file to ensure workstations do not spend time competing for record locks. Even with 22 workstations active, all the private files are comfortably cached by the file server, so this test also has a minimal impact on the file server's disks.

We ran each test using the standard database, containing mainly space characters in each 30 Kilobyte record. We also used two variant databases with records of the same size but filled with consecutive repeats of one or other of the two data patterns expected to cause either maximum or sub-maximal baseline wander. Any level of baseline wander was expected to reduce the total data throughput of the 100-BaseTX network, but it was thought to be possible that the maximum level might actually have a less serious effect on data throughput because it could force the 100-BaseTX interface to reset periodically and return to full-speed operation for a short time, whereas the lower level might cause continuous poor performance without forcing the interface to reset.

Since we were being asked to compare the 100-BaseT and 100VG technologies rather than different implementations of those technologies, we ran the tests using 3 different EISA 100-BaseT network interface cards—from 3Com, Intel, and SMC—in the file servers, and various combinations of 100-BaseT hubs from NetWorth, SMC, Bay, and 3Com. All the 100VG kit was from HP. The ISA workstations were fitted with Microdyne 100-BaseT or HP 100VG cards.

Multiple servers and clients

The main function of our network performance tests is to put an enormous load on the networks under investigation. In the test series reported here, we were concerned to ensure that the 100VG and 100-BaseT networks were fully saturated and carrying the maximum possible amount of data. We could then be certain that measurable differences in their performance were due to real differences between the technologies and were less likely to reflect the quality of the currently available hubs and workstation and server network interface cards.

It was clear from preliminary tests that a single NetWare file server could not be guaranteed to saturate either network technology. In contrast, test networks constructed with two or three file servers, in which approximately equal numbers of clients communicated with each server (one server per client) showed clear signs that the communications channel was saturated:

- Maximum throughput with 3 servers is not greater than with 2.
- With 2 or 3 servers, each server operates at a substantially lower rate than the maximum recorded for a single server.
- Maximum throughput is usually reached with 9 or 10 active clients and thereafter does not substantially increase with increasing numbers of clients.

We have therefore used the results from networks with two and three servers to draw our main conclusions about the total throughput possible with each technology. These results lead us to three major observations:

- 100-BaseT and 100VG are both capable of moving data in the direction server-to-client at the maximum expected rate.
- 100VG can sustain bi-directional data flows at the maximum expected rate, but 100-BaseT is always slower.
- Data patterns designed from a theoretical consideration of the 100-BaseTX specification to cause baseline wander are transmitted at a substantially reduced rate in many network configurations compared with normal data.

We calculated the maximum expected throughput for our tests on 100 Megabits per second networks by scaling up results obtained over many years with conventional 10 Megabits per second Ethernet. On 10BaseT Ethernet our Overlaid Read and Read/Write tests run at a maximum speed of about 1080 Kilobytes per second (8.64 Megabits per second). The difference between 8.64 and 10 Megabits per second—the actual channel speed for Ethernet—represents the overhead introduced by the encapsulation of the raw data in IPX packets and the detailed operation of NetWare's IPX protocol. This overhead would be expected to remain exactly the same as the speed of the network is scaled up, given the same packet format, so the maximum throughput for a 100 Megabit per second network would be expected to be 86.4 Megabits per second.

Indeed we found that both 100-BaseT and 100VG almost achieved this with a maximum throughput on the Overlaid Read test of about 10700 Kilobytes per second (85.6 Megabits per second). In the Overlaid Read test, data flows exclusively from server to client, although protocol control packets must also pass from client to server. All 100VG configurations with 2 or three servers easily achieved the theoretical maximum throughput whereas only one of the many 100-BaseT configurations—that with 2 servers and Intel server cards—could reach it. Graph 1 on page 22 shows the typical curve of throughput against number of active clients in a network with 2 servers, in which the maximum 100-BaseT throughput was a little lower than the maximum for 100VG. The curves for 100VG with all three patterns are identical, whereas both patterns 1 and 2 reduced 100-BaseT throughput to less than 50% of that seen with the normal pattern.

Performance of 100-BaseTX and 100VG (*continued*)

100VG also achieved almost the same maximum throughput in the Overlaid Read/Write test, in which equal amounts of data flow between client and server, and it was at least 26% faster than any 100-BaseT configuration in that test (Graph 2, page 23). Again, all three data patterns gave an identical throughput with 100VG, and in this case only pattern 1 had any effect on 100-BaseT throughput, causing a lesser reduction than that seen in the Overlaid Read test.

Single client

In our standard lab test for network throughput, multiple clients simultaneously access a small number of file servers. We also investigated the performance of 100VG and 100-BaseT in a completely different environment, in which a single fast client was connected to a single server. Graph 3 on page 24 shows the resulting throughput for every possible combination of 80486 66MHz EISA client and server cards. There are wide variations between the different 100-BaseT cards; the combination of two 3Com cards was the fastest at 7736 Kilobytes per second (62 Megabits per second) and the combination of two Intel cards the slowest at 4259 Kilobytes per second (34 Megabits per second). 100VG ran only slightly faster than the slowest 100-BaseT configuration. The large variation in 100-BaseT performance with different interface cards presumably reflects differences in hardware and software design philosophies between the currently available cards. Since the Intel card was the slowest with a single client but the fastest with multiple clients, it seems likely that different cards are optimised for different network configurations.

As with the multiple client networks, the Overlaid Read/Write throughput for all the 100-BaseT configurations with a single fast client was markedly reduced compared with the Overlaid Read figure whereas the 100VG single client throughput was hardly changed (Graph 3).

Patterns 1 and 2 reduced 100-BaseT throughput in the single-client Overlaid Read tests, but to a lesser extent than with multiple clients.

Multiple 100-BaseT hubs

In order to be absolutely certain that our 100-BaseT results with patterns 1 and 2 did not just reflect problems with one particular interface card or hub, we constructed networks with many configurations of two 100-BaseT hubs and three servers, each with one of the three server 100-BaseT cards. These included all 5 possible combinations of two makes of hub, given the lack of an uplink port on the SMC and 3Com hubs. All the 100-BaseT devices operated successfully together with the standard data patterns, and they all showed between 20% and 80% performance degradation with pattern 1 and usually slightly less degradation with pattern 2.

Cable length

Theoretical considerations of baseline wander in 100-BaseTX suggest that the problem should become more severe with increasing cable length, even within the limits imposed by adherence to the category 5 specification. Our main series of tests was carried out using 97 metre Category 5 certified cables throughout in an attempt to reproduce the worst-possible legal case for 100-BaseTX. We also investigated the effect on throughput, with and without baseline wander, when the cables between the hubs and the servers, the clients, or both simultaneously, were shortened to 10 metres. Cable length had no effect on throughput with the standard pattern and an unpredictable effect on throughput with baseline wander. Shorter cables did not always produce greater throughput than longer ones, and 10 metre cables all round did not eliminate the baseline wander effect.

Complete communications failures

Probably the only safe conclusion from the tests on varying cable lengths is that the effect of baseline wander on throughput is highly susceptible to relatively minor changes in the physical medium. Throughput with 100VG was never affected by changes in cable length.

We observed several complete failures of transmission in which the client(s) lost contact with the server(s), exclusively with the Overlaid Read test and pattern 2. These were consistent, in that every repeat of the test with the same configuration failed in the same way, and were seen with nearly all configurations that included a 100-BaseT Bay Networks hub or an Intel card in a client, including the single server and client networks. When there was such a failure, transmission could only be reactivated by power cycling the Bay hub or reloading the Intel server network drivers.

Variability in Measured Throughput

We report the Standard Error in the Mean for all our results, calculated as the sum of the Standard Errors for each individual client participating in a test. This gives an indication of the variation in throughput that was experienced by the clients when performing successive repeats of the same test activity; a higher standard error means greater variability. In general, 100-BaseT was much more variable than 100VG in all network configurations, particularly in the Overlaid Read/Write tests, where the standard error for 100-BaseT often rose to over 5% of the total throughput. For 100VG it was typically below 0.5%.

Conclusions

The Netmarq performance tests are intended to measure the performance of network components and technologies under conditions that are as near real-life as possible and are therefore directly relevant to network users. We therefore wrote our tests in Clipper, an end-user database programming language, and we measure real data throughput between multiple clients and servers. We strongly believe that the actual data throughput between server and client that our tests measure is a better indication of real network performance than other more artificial measures such as packets forwarded per second; network users expect to move data from one place to another, not raw packets.

By using a range of 100-BaseT equipment from different manufacturers we tried to ensure that our results reflected the general performance of the 100-BaseT technology, rather than faults in individual products. In fact our use of only the HP 100VG interface cards and hubs would be expected to result in a bias against 100VG. Nevertheless, 100VG proved a faster and more stable network technology than 100-BaseT in the majority of our tests. Although both 100VG and 100-BaseT were capable of moving data from server to clients at the expected maximum rate, only 100VG achieved this rate when simultaneously moving data both from server to clients and from clients to server. In addition, 100-BaseT networks proved highly susceptible to data patterns designed to cause baseline wander, to the extent of complete communications failure in a few cases.

The only test in which 100-BaseT was possibly superior to 100VG was that with a single client and server and uni-directional data flow—most (but not all) 100-BaseT network interface cards gave a faster network than 100VG. With bi-directional data flow both the technologies operated at substantially the same speed in this configuration. Given the great variation between different 100-BaseT cards, we cannot be certain that this is a real difference between the technologies until we have tested a similarly wide range of 100VG cards.


ROGER COHEN holds a science degree from Cambridge University and spent 4 years doing postgraduate research in medical biochemistry. After careers in lexicography and text processing, he became network manager for various major UK companies. Since 1989 he has been continuously developing a test suite for assessing the performance of network components, technologies and operating systems, as well as designing and troubleshooting large LANs and WANs. He is now technical director of Netmarq, a specialist network test house, and also of Homepage, a Web design and service provision company. E-mail: netmarq@cix.compulink.co.uk

continued on next page

Performance of 100-BaseTX and 100VG (continued)


Graph 1

Overlaid Read test with 2 file servers, 3Com 100BT server cards and NetWorth hub



Graph 2


Overlaid Read/Write test with 2 file servers, 3Com 100BT server cards and NetWorth hub



Performance of 100-BaseTX and 100VG (continued)

Graph 3

Data Throughput with Single Client and Server



TERENA, the Trans-European Research and Education Networking Association

by Karel Vietsch, TERENA

Abstract

TERENA, the *Trans-European Research and Education Networking Association*, was established in October 1994 through the merger of the RARE [1] and EARN associations. This article presents an overview of the current activities of TERENA, its role in the European networking environment and its plans for the near future.

Status of networking in Europe

Perhaps the most significant development in networking in the mid 1990s is the fact that companies and private individuals are gaining access to the Internet in large numbers. Commercial companies are entering the market to provide connectivity and information services to these new users: Internet Service Providers, mostly small and new companies, offer access to the Internet, and more and more diverse sources of information are made available via the Internet on a commercial basis. To a very large extent these new access and information services are based on technologies and methodologies that were developed and introduced in the research networking environment during the past ten years.

At the same time this important development does little to fulfill the present needs of network users in research establishments and institutes for higher education in Europe. The services they need and which require increasingly pervasive high-bandwidth network access, are not offered by commercial access providers, Public Network Operators or other companies. Internet Service Providers have little to offer to users in the European research community, except in regions where because of infrastructural or financial limitations no alternatives are available. National research network operators—dedicated and professionally managed organisations which since the mid 1980s have been established in all countries in Europe—remain indispensable to provide to the research community the more advanced services it needs. At least for the research community, networking is by no means a commodity yet.

At present the European research networking organisations are faced by various problems. In many parts of Central and Eastern Europe the limitations of the available underlying telecommunications infrastructure continue to pose severe restrictions on the development of research networking. Here considerable investments are needed. Although research networks in Central and Eastern Europe can point at significant achievements, it remains difficult to catch up with the developments in Western Europe and North-America.

At the same time those networking organisations which are ready to introduce more advanced and high-capacity services, find themselves limited by financial constraints. The prices that telecommunications operators charge for international and intercontinental high-speed connections are extremely high, not related to real costs and to a large extent beyond the financial means of the research community. Consequently there is a danger that the national electronic highways for the research community that are being established in many countries in Europe, North-America and the Pacific Rim, will be interconnected by electronic cart tracks only. It is a major challenge for the next few years to bring down the prices of high-speed connections and to convince telecommunications operators that they should see research networking not as a money spinner but as a partner for strategic collaboration in the development of new technologies and services.

continued on next page

TERENA (continued)

TERENA

In this situation, where research networking organisations in Europe have many joint interests and face many common problems, there is an obvious need for them to work together. TERENA, the association that almost all national research networks in Europe belong to, is the vehicle to organize this collaboration.

TERENA has *National Members*—the organisations responsible for the management of the national research network in their country—*International Members*—international treaty organisations with a strong interest in networking for research—and *Associate Members*. At the end of 1996 forty-one countries were represented in TERENA.

National Members:

INIMA (Albania)	IIUCC (Israel)
ACOnet (Austria)	CNUCE (Italy)
Academy of Sciences (Azerbaijan)	LITNET (Lithuania)
UNIBEL (Belarus)	RESTENA (Luxembourg)
BELNET (Belgium)	MARNET (FYRoMacedonia)
UNICOM-B (Bulgaria)	Cntr. for Informatics (Moldova)
CARNet (Croatia)	SURFnet (Netherlands)
CYNET (Cyprus)	UNINETT (Norway)
CESNET (Czech Republic)	NASK (Poland)
UNI-C (Denmark)	FCCN (Portugal)
FRCU (Egypt)	Min. Res. & Techn. (Romania)
EENet (Estonia)	RFBR (Russian Federation)
FUNET (Finland)	SANET (Slovakia)
RENATER (France)	ARNES (Slovenia)
Academy of Sciences (Georgia)	RedIRIS (Spain)
DFN (Germany)	SUNET (Sweden)
GR-NET (Greece)	SWITCH (Switzerland)
HUNGARNET (Hungary)	TUVAKA/ULAKBIM (Turkey)
SURIS (Iceland)	Academy of Sciences (Ukraine)
IPM (Iran)	UKERNA (United Kingdom)
HEAnet (Ireland)	

International Members:

CERN	ECMWF
------	-------

Associate Members:

CEENet	ESA
DANTE	EUnet CS B.V.
Digital Equipment Corp.	JINR
EMBL	NORDUNET

Table 1: TERENA Membership (December 1996)

TERENA's functions

The mission of TERENA is the promotion and development of a high-quality international information and telecommunications infrastructure to support European research and education. This translates into four categories of activities, that form the pillars of the association:

- To represent the interests and the opinions of TERENA's members in contacts with governments, funding bodies, telecommunications operators, industries and the public at large—both at the European level and in the relations with similar organisations from other continents;

- To carry out a Technical Programme where technical experts from the TERENA member countries work together in Working Groups, Task Forces and innovative and collaborative Projects, to further develop and promote networking technologies and services for the research community;
- To organize conferences, workshops and seminars to exchange information between the TERENA members and to make them and others involved in the wider networking environment aware of relevant developments;
- To provide Europe-wide utility-type services on a neutral basis to the TERENA community and, if needed, to the networking community at large as well.

In the last year TERENA has seen a sizeable increase in the activity level. Staffing problems have been solved, and a complete technical and administrative staff is now available to support TERENA's activities. The Technical Programme is gaining impetus and new Task Forces and Projects are being started. TERENA activities are growing in number, size, participation and speed. Thus the association is in a good position to answer the challenges that are posed by the present status of research networking in Europe.

TERENA Executive Committee:

Stefano Trumpy	President
Steve Druck	Vice President Services
Brian Gilmore	Vice President Technical Programme
Wulf Bauerfeld	Vice President Conferences
Lajos Bálint	Treasurer

TERENA Staff Management:

Karel Vietsch	Secretary General
John Dyer	Chief Technical Officer
Fay Howard	Executive Officer
Daniel Karrenberg	RIPE NCC Manager

Table 2: TERENA Management (December 1996)

Technical Programme

TERENA's Technical Programme is organized through Working Groups, Task Forces and Projects.

The *Working Groups* structure stems from the RARE Working Groups that were set up in the late 1980s [1]. Currently TERENA has seven Working Groups, for the areas of Internationalization, Information Services and User Support, Lower Layer Technologies, Mail and Messaging, Networked Multimedia Applications, Quality Management for Networking, and Security Technology. The Working Groups form a huge pan-European resource in the form of many e-mail distribution lists reaching hundreds of people working in the computer communications field. The Groups usually meet at TERENA's annual conference. Each Working Group has a convenor, and all convenors together with a few invited well-known experts form the *TERENA Technical Committee* (TTC) under the chairmanship of TERENA's Vice President Technical Programme.

TERENA (continued)

Task Forces are small groups that work on well-defined and bounded issues. Their life time may vary from months to a few years. Task Forces are established by the TERENA Technical Committee, and they are reviewed at least annually by the TTC, at which point their mandate to continue is assessed. Task Force activities are expected to be self funded by the members making their own arrangements for travel and subsistence. Task Force meetings usually take place at the TERENA Secretariat in Amsterdam.

Projects are activities that are undertaken by a small subset of either Working Group or Task Force members—or indeed by any other small group of people—working on behalf of TERENA. Projects are governed by a single *Controlling Document* defining justification, objectives, deliverables and resource requirements of the project, as well as financial commitments to the project and explaining the change control mechanism for the project. Projects will fall into one of three categories. “Minor projects” take on a small amount of work that can be financed from TERENA resources, without the need for specific additional fund raising. “Pilot projects” are small preliminary explorations of a topic or solution to be funded from TERENA’s Pilot Project Fund; such activities are expected to lead to a major community-wide project. “Major projects” are large projects, normally following from a successful Pilot project, and funded by voluntary subscription to particular projects by TERENA members and/or by subsidies from third parties.

Brian Gilmore	Chairman
John Dyer	Secretary
Claudio Allocchio	WG-MSG
Rob Blokzijl	
Manfred Bogen	WG-QMN
Rudiger Grimm	WG-SEC
Dave Hartland	WG-ISUS
Borka Jerman-Blazic	WG-I18N
Olav Kvittem	WG-LLT
John Kwan	WG-NMA
Olivier Martin	
Karel Vietsch	TERENA Secretary General

Table 3: TERENA Technical Committee (December 1996)

Some recent developments in the Technical Programme are the following:

- TERENA and the Université Libre de Bruxelles are the project partners in the SCIMITAR project, a concertation project for the “Telematics for Research” sector of the 4th Framework Programme of the European Union;
- TERENA, SURFnet and UKERNA released a CD-ROM containing documentation and code for the latest versions of PGP (*Pretty Good Privacy*); this CD-ROM is being distributed at several international conferences and additional copies may be ordered from [<pgp-cdrom@terena.nl>](mailto:pgp-cdrom@terena.nl);
- Early 1996 saw the start of the DEVICE project, which undertakes a comprehensive survey of desktop video conferencing products and interoperability testing on selected products;

- A Task Force, called TF-TEN, has been set up to co-ordinate—as part of the TEN-34 project—the involvement of the research community in the European broadband ATM network project;
- A contract has been placed for the first phase of the project to update and enhance the “Guide to Network Resource Tools”;
- A project has been started to update and re-publish a technical report on multimedia expertise and projects;
- Preparations are made for a pilot for European Security Incident Response Coordination;
- A Task Force, called *TF-Cache*, has been established to promote the deployment of caching hierarchies in Europe, to link these with initiatives in the United States and to provide information, documentation and assistance for those wishing to start a caching service;
- The new *TF-ETINU* will identify, discuss and recommend methods of providing on line services and support in such a way that the user can take full advantage of these, using readily available browsers;
- A new Task Force, *TF-ETM*, has been set up to prepare for a project on network metrics and statistics collection.

Conferences

Annually TERENA organizes a large international conference, known as JENC—the *Joint European Networking Conference*. The JENC7 conference, titled “Networking in the Information Society,” took place in Budapest on 13–16 May 1996. The conference, well attended with over 320 participants from over 38 countries world-wide, was judged as very successful. Topics of 46 papers ranged from ATM technology issues and network programming models via applied security technology and tele-teaching to in-depth discussion on the economics of the Internet. The “Telematics for Research” sector programme was also well covered during the conference. All plenary sessions were broadcast on the MBone, using the infrastructure of the German ATM pilot.

JENC8 will take place in Edinburgh, Scotland on 12–15 May 1997. The theme will be “Diversity and Integration: The New European Networking Landscape.” The year 1997 marks the final stage of deregulation of much of the European telecommunications infrastructure. This will have a great impact on the European networking community. JENC8 will focus on the effects of this change for both the academic and commercial worlds; addressing also the strengthening of the European position in the competitive, global environment. In addition the conference will explore state of the art network technology issues from a technical and users’ viewpoint.

In the week preceding the JENC8 conference TERENA will organize in Edinburgh an Advanced Networking Workshop, with financial support from NATO. The theme of the workshop will be “Migrating Towards a High Speed Networking Service.” Networking in Central and Eastern European countries is developing at a very rapid rate. The workshop is aimed at those who are currently running national network services but who will, in the near future, be moving to new technologies to provide new applications, higher bandwidth and a more reliable service.

TERENA (*continued*)

The Edinburgh workshop stands in a long tradition of TERENA (formerly: RARE) organized and NATO funded events to support the development of networking in Central and Eastern Europe. Earlier workshops took place in Prague (1994), Warsaw (1995), and Budapest (1993–94, 1995 and 1996).

Plans for other TERENA workshops in 1997 are still under consideration. Possible topics are the policy and managerial aspects of running a national research networking organization, and the development of networking in the Mediterranean area.

GUM NCC service

In May 1995 the GUM NCC (*Global Updating & Monitoring Network Coordination Centre*) was created to continue the co-ordination of NJE activities. Under a contract with TERENA the GUM NCC office at Nijmegen University, the Netherlands, takes care of the management of the existing NJE infrastructure by providing assistance on the daily operation to the core sites, co-ordination of the backbone and the international topology, management of BITEARN NODES, production of routing tables and production of traffic reports, on a world-wide basis. While there is a clear trend that NJE traffic is becoming relatively less and less important for the research networking community and the number of NJE nodes is decreasing, there is still a considerable amount of NJE traffic. Mid 1996 the NJE infrastructure contained 8 core sites in Europe serving a total of 28 countries.

Some of these 28 countries have only a low-speed unreliable connection to the rest of world-wide networking. Therefore a working NJE connection was the only feasible solution for them. The problem of reliability and capacity of connections is no doubt the most important issue determining the dependence of less-advanced countries on NJE traffic. Other obstacles for converting to other transmission protocols may concern organisational and political issues. Experience has shown that even in more advanced countries an orderly winding down of NJE traffic takes some effort and time.

Therefore TERENA decided that the GUM NCC service will be wound down gradually. No part of the service will continue after 1997. All users of the GUM NCC service are urged to remove any remaining dependence on NJE traffic as soon as possible; the GUM NCC office at Nijmegen will be available to provide advice and assistance on converting to other transmission protocols.

RIPE NCC service

TERENA's RIPE NCC (*Réseaux Internet Protocol Européens Network Coordination Centre*) service has been described in this journal some three years ago [2]. Since then the size and scope of the service has seen a tremendous development. At present, the RIPE NCC activities can be grouped into four main categories:

- *Registration services*: This represents activities related to RIPE NCC's role as *Regional Internet Registry for Europe* and the surrounding areas. It includes handling of requests for assignment or allocation of IP address space, management of reverse domains associated with this address space as well as auditing and quality control to ensure fair and expedient processing of requests. Also included in this area are training of *Local Internet Registries*, production of documentation related to Internet registration and specific activities to ensure a proper and expedient start of new Local Internet Registries. Services performed in this area are only accessible to formally established Local Internet Registries contributing to the funding of the RIPE NCC.

- *Co-ordination activities:* The activities have as a common purpose to support the coherent operation of the Internet in the European area. An important activity is the provision of access to the RIPE database which provides information about address space and routing policies together with the appropriate contact points. Developing and publishing the RIPE database software is also part of this area, as is the provision of information services for Internet Service Providers and the general public via the Internet. Operational co-ordination such as efforts to reduce the number of globally visible routing prefixes also falls into this category, as does the production and publication of software tools for such efforts. In order to be effective the services in this area have to be accessible to the general Internet public. Contributors to the funding of the RIPE NCC receive precedence over other users when special support is needed.
- *Administration activities:* This area covers all regular reports published by the RIPE NCC, administrative support for RIPE as well as general administrative overheads. As such it includes production of the Quarterly Reports and the resources needed for charging, billing and the general financial administration.
- *New activities:* This area represents those activities that cannot be fully specified at the time of budget preparation. The existence of this area gives the RIPE NCC the flexibility to react quickly to the rapid changing needs in today's Internet.

The number of Local Internet Registries, the prime users of the RIPE NCC services, has grown dramatically over the past years, to about 500 by the end of 1996. A further growth to almost 1000 Local Internet Registries by the end of 1997 is foreseen. Although economies of scale and new, automated working methods greatly help the efficiency of the RIPE NCC operations, of course this tremendous growth in the size of the operation has resulted in a substantial growth of the—TERENA employed—RIPE NCC staff. Having started with three people in 1992, the RIPE NCC staff now counts 17 employees, and a further growth to 27 people is expected by the end of 1997.

In view of this growth, possible new forms of the legal structure for the RIPE NCC service are being investigated, with the objective to find and implement a new legal status which will provide a stable basis and limit the mutual liabilities between RIPE NCC and the other current TERENA activities.

TERENA's challenges for 1997

TERENA's activities in 1997 largely follow from the work that has been done in recent years. The activities in the Technical Programme will be increased and intensified, the scope of the conferences will be somewhat widened, and a significant growth is expected for the RIPE NCC service. However, new emphasis will be put on certain activities, in order to answer some of the challenges that are facing the TERENA membership.

A first challenge is to better fulfill TERENA's role as a spokesman and a representative for the interests of its members. Many of the problems that research networks have to deal with today, are related to financial, regulatory and policy issues at a European level. It is therefore very important to make the political case for research networking, to point out the indispensable role of advanced network facilities for scientific research and education, and to underline the distinguished role of the research community in establishing overall technical advancement in the area of networking.

TERENA (*continued*)

Of course this requires investments on TERENA's part. The urgent need to improve the relations between the European research networking community and governments, funding agencies, telecom operators and industries, makes these investments well justified.

A second challenge is to enable TERENA's members to benefit from its added value as a mechanism to promote communication, the exchange of information and collaboration between networking organizations. The key word here is "interactivity," and that objective will receive more attention in all of TERENA's activities, in particular in the representation area, the Technical Programme and the conferences.

A particular example of promoting communication, exchange of information and collaboration between networking organizations is the transfer of knowledge to technologically developing countries. In the recent past, TERENA has already been active in this field, in particular towards countries in Central and Eastern Europe and the area of the former Soviet Union. That line of activity will be continued, and at the same time these activities will be extended to other regions, in particular the Mediterranean area. Funding for such activities will be sought from international organizations and from the relevant countries themselves, while TERENA will be willing to invest manpower from its existing staff to develop this line of activities.

Another question is related to the start-up of new projects in the Technical Programme. The principle here will remain that these projects are expected to be self-financing, and therefore additional contributions are asked from TERENA's members to fund the projects. However it is understandable that networking organisations are hesitant to invest money in projects which have little to show but a plan on paper. Therefore, as an experiment, a fund—the Pilot Project Fund mentioned above—has been created from TERENA's own resources to provide "pump priming" money to new projects. In this way projects can make a start with little or no additional funds being collected, and it is expected that this will make it easier to find funding for these projects once they have shown that they represent a feasible plan.

Finally, a major challenge is to execute all these enlarged and intensified TERENA activities in 1997 without an increase of the membership fees. TERENA feels it should not ask its members for an increase of the membership fees without showing first that it is capable to make a success of its ambitious plans and to provide a return on investment to its members. Therefore the year 1997 will aim to provide to the TERENA members "more value for the same amount of money."

Further information

The material in this article has been gathered from a number of publicly available TERENA documents. It is impossible to give a full detailed overview of all of TERENA's activities here. Also, developments in our organization are so fast that any printed information runs the danger of being outdated very soon. The interested reader is therefore referred to the complete and up-to-date information that is available on the TERENA Web site:

<http://www.terena.nl>

The RIPE NCC service has its own Web site at:

<http://www.ripe.net>

References

- [1] "Réseaux Associés pour la Recherche Européenne," *ConneXions*, Volume 6, No. 1, January 1992.
- [2] Karrenberg, Daniel, "The RIPE NCC and the Routing Registry for Europe," *ConneXions*, Volume 7, No. 11, November 1993.
- [3] *ConneXions*, Volume 7, No. 5, May 1993, "Special Issue: Focus on Europe."
- [4] Crowcroft, J., "A Brief Glossary/Overview of the European Internet," *ConneXions*, Volume 9, No. 12, December 1995.
- [5] Stockman, B., "Current Status on Networking in Europe," *ConneXions*, Volume 5, No. 7, July 1991.
- [6] Stockman, B., "EBONE, The European Internet Backbone," *ConneXions*, Volume 7, No. 5, May 1993.
- [7] Stockman, B., "Global Connectivity: The Global Internet Exchange (GIX)," *ConneXions*, Volume 7, No. 11, November 1993.
- [8] Bersee, J., "Profile: DANTE and EuropaNET," *ConneXions*, Volume 8, No. 6, June 1994.
- [9] Kowack, G., "Profile: EUnet," *ConneXions*, Volume 7, No. 11, November 1993.
- [10] Greisen, Frode, "Ebone and other parts of the European Internet," *ConneXions*, Volume 10, No. 8, August 1996.
- [11] European Round Table of Industrialists: "Building the Information Highways," June 1994.
- [12] Crowcroft, J., "International Internetworking," *ConneXions*, Volume 2, No. 4, April 1988.
- [13] Goldstein, S. & Michau, C., "Convergence of European and North American Research and Academic Networking," *ConneXions*, Volume 5, No. 4, April 1991.
- [14] Guy Almes, Peter Ford, and Peter Lothberg, "Proposal for Global Internet Connectivity," June 1992.
- [15] Tony Bates, Daniel Karrenberg, Peter Löthberg, Bernhard Stockman and Marten Terpstra, "Internet Routing in a Multi Provider, Multi Path Open Environment," February 1993.
- [16] Tony Bates, Jean-Michel Jouanigot, Daniel Karrenberg, Peter Lothberg and Marten Terpstra, "Representation of IP Routing Policies in the RIPE Database," March 1993.

KAREL VIETSCH holds an M.Sc. and a Ph.D. from Leiden University, the Netherlands. He was a teaching and research assistant at Leiden University (1973–1979) and General Manager of the Department of Mathematics and Computer Science of Delft University of Technology (1980–1984). He joined the Dutch Ministry of Education and Science in 1984 as a Senior Policy Advisor for Information Technology in the Science Policy Division. From 1992 until 1996 Dr. Vietsch was Head of Unit for Information and Infrastructure in the Research and Science Policy Directorate. In these functions he was responsible for the government involvement in the establishment and funding of the Dutch national research network SURFnet. Dr. Vietsch represented his government in European collaborative actions for research networking, notably the Eureka projects COSINE (1985–1993) and EuroCAIRN (1994–1996) and the current European Networking Policy Group (ENPG), of which he was a co-founder. He took up the duties of Secretary General of TERENA in March 1996. E-mail: vietsch@terena.nl

SNMPng Advisory Team Status Report

Abstract

At the Montreal IETF, the IETF Network Management Area Director announced the intention to form an advisory team to analyze the proposed, existing approaches for SNMP security. These approaches are commonly known as USEC and v2*. (See the May 1996 issue of *ConneXions* for further details on each approach). The general direction to the advisory group was to analyze the two approaches and provide recommendations to the Network Management community. This article provides a brief background on the activities and the status of work (as of December 1996) including a very high-level description of the recommendations.

Background

During the past several years, there have been a number of activities aimed at incorporating security improvements to the *Simple Network Management Protocol* (SNMP). This was one of the principal purposes of developing version 2 of SNMP. Unfortunately, strongly held differences on how to incorporate security into SNMP prevented the SNMPv2 Working Group from coming to closure on a single security approach. As a result, two approaches have emerged. These approaches are commonly known as USEC and v2*.

Advisory Team

At the 36th IETF in Montreal, the Network Management Area Director announced the intention to form an advisory team to analyze the two proposed approaches for SNMP security. The chartered name for the group is the *Security and Administrative Framework Evolution for SNMP Advisory Team*. Since the group has mostly referred to itself as just the "Advisory Team," that terminology will be used for the remainder of this article. The group is made up of the following individuals:

David Harrington
Jeff Johnson
David Levi
John Linn
Russ Mundy (Chair)
Shawn Routhier
Glenn Waters
Bert Wijnen

As stated in the charter, the Advisory Team intended to publish one status report at the end of September and a white paper in early November. This would provide information for the network management community a month prior to the December 1996 IETF in San Jose. As with most IETF activities, the Advisory Team is supported essentially on a volunteer basis and progress was not as rapid as anticipated. Short but not particularly informative reports were provided in early October and November.

In accord with the Team's Charter, a large portion of the work was done via a private e-mail list. This resulted in steady but slow progress in resolving differences between the two approaches and developing recommendations. In November, the Team decided that we were not satisfied with our progress on the e-mail list and that a face-to-face meeting was needed. As a result of this meeting, we resolved all of the larger issues necessary to define a merged, single approach. Though some areas of disagreement still exist, we do not believe that these are critical to the overall recommended approach. This article provides a very high-level description of the recommended approach from the Team. We provided a presentation during the recent IETF meeting, and the slides as well as additional documents are available at: <http://www.tis.com/docs/research/network/snmp-ng.html>

Advisory Team process

As described in the Charter, our process was to identify the commonalities of the proposals that can be merged and the differences between the proposals including understanding the requirements that drove these differences. Based on these commonalities, differences and requirements, a set of recommendations and rationale were to be developed.

The Team reviewed material from mailing lists as well as USEC and v2* publications to identify areas of agreement and disagreement. We also discovered that there were some areas that we were initially unsure about whether there was agreement or disagreement between the approaches. The Team discussed and debated the various areas that were identified and categorized them into a set of issues to better focus our discussions. The names of the initial issues as posted to the SNMP and SNMPv2 mail lists on November 8th are:

- Timeliness Check Module
- Authentication Module
- Encryption Module
- Proxy Determination Module
- Proxy Handling Module
- Access Control Module
- Varbind Processing Module
- Overall Framework Issues
- The Layers Issue

As we worked through our process, each of these issues had a listing of items on which USEC and v2* agreed and items on which there was disagreement. There was also a list of items that the Team was unsure about. The Team was able to reach agreement and develop recommendations on a number of items through the e-mail list but when November arrived, the Team was not satisfied with our progress. Consequently, we met to work out our differences on unresolved areas and confirm consensus in areas where there was less than full agreement. During the face-to-face meeting, there was a significant amount of discussion about how we should categorize the various issues and items. We worked rigorously to understand areas and causes of disagreement as well as requirements driving implementation choices. The result of our efforts is definition of a set of modules, sub-modules, structures and interfaces for SNMP messages and processing.

In both the e-mail list and meeting forums, we had good technical exchanges and debates including disagreements on SNMP philosophies and approaches but at no time in this process did any of the exchanges become acrimonious. As a result of these exchanges, we were able to define an approach which incorporates pieces and concepts from both USEC and v2*. Although the approach defined by the Team will require some changes to current USEC and v2* implementations, we believe this approach should result in the merger of USEC and v2* back into a single standard. The approach defined by the Team essentially constitutes our recommendations. The approach will be described at a high level in the following section of this article and additional detail was provided at the December 1996 IETF

**The recommended
approach:
A High Level View**

The Team considered using either USEC or v2* as the basis for our recommendations but rejected that approach. We concluded that it would be more effective to use portions of both approaches in developing our recommendations rather than trying to modify either the USEC or v2* approach.

continued on next page

SNMPng Advisory Team Status Report (*continued*)

The Team's recommended approach defines a set of modules, sub-modules and interfaces for an SNMP engine. It identifies interfaces to associated applications that can perform functions external to the SNMP engine. The approach also defines revisions to a part of the message structure.

The Administrative and Security Framework of the recommended approach are built to the greatest extent possible on the existing base. For example, the PDU portion of the message structure comes from RFC 1905. It was our general consensus that the user based security framework from the USEC approach was acceptable for the v2* approach and, therefore, provided a large amount of the detail for the security specific portions of our recommended approach. Additionally, we believe that the recommended approach is generally compliant with the current SNMP Standards Track RFCs (RFC 1902–1908).

The requirements, functions and modules were carefully examined by the Team. During the process, we were able to define three modules (some with sub-modules) that perform the principal functions of an SNMP engine. The current names for these modules are *Message Processing and Control*, *Security Model*, and *Local Processing*. The Message Processing and Control module handles SNMP message creation and parsing functions. In some ways it can be thought of as the “traffic cop” for the SNMP engine including determining if proxy handling is required for any particular SNMP message. The Security Model module provides authentication and encryption functions. This module also checks the timeliness of certain SNMP messages. The Local Processing module performs access control for varbind data, processing varbind data and trap processing.

The interfaces between these modules and applications for the principal SNMP functions, such as generating or receiving a request message, have been defined. In general, we believe that modules, sub-modules and applications can be replaced independently without affecting other components provided that the new components maintain the same interface. We believe the recommended approach permits proxy functions as well as network management station functions to be placed in applications with defined interfaces rather than being entwined with the protocol engine. Another portion of the recommended approach identifies the data required in the message header.

To help define the recommended approach, we have developed illustrations for the message structure and scenario diagrams. The message structure illustrations show the header information needed for this modular approach while the scenario diagrams identify the sequence of events and the information that must pass between the modules and/or application. We are currently in the process of cleaning up the drafts of these illustrations and have begun documenting the textual descriptions of the events in the scenario diagrams.

Additional material

We are continuing to work on the details of our recommended approach. Also, we plan on making additional material, such as scenario diagrams, available via the web. When this material becomes available, it will be announced on both the SNMP and SNMPv2 mail lists.

Conclusion

Our recommended approach should be considered work in progress and not attacked because it is incomplete. We hope our approach will provide a sufficient basis for transfer to a Working Group that is chartered to carry this merger forward to a single network management standard for the IETF community.

—Russ Mundy (mundy@tis.com)

Book Review

Network Security: Private Communication in a Public World, by Charlie Kaufman, Radia Perlman, and Mike Speciner, ISBN 0-13-061466-1, Prentice Hall, 1995.

The number of technical textbooks that are comprehensive, sound, and have an attitude must be minuscule, and I'd not expect to find one in the field of computer security, traditionally a field desiccated by its own pontification. *Network Security* is a refreshing departure from the norm.

Attitude

Forged in the crucible of the MIT hacker community, when "hacker" was a term of honor, the authors guide the reader through intricacies of network security technology with a deft hand, one that flashes the famed MIT brass beaver flagrantly. A casual attitude combined with dead-on technical mastery pervades the writing.

Network Security has enough "meat" to be enjoyed by engineers and students. It is a good reference, and the problems at the end of each chapter make it suitable for teaching. The book manages to include much detailed information and good qualitative explanations.

Organization

The first third of the book covers cryptographic algorithms. This is standard material, essential to the rest of the book. The number theoretic algorithms are explained and bolstered with short discussions of their mathematical underpinnings. There are a few surprises: the smooth number weakness of RSA, prime generation techniques for Diffie-Hellman, and zero-knowledge schemes.

Lucid

Protocols for authentication in a networked environment constitute the second part of the book, and this is the strong suit of *Network Security*. The chapters take the reader through an increasingly sophisticated series of protocols for security functions (authentication, privacy, integrity), clearly illustrated. By describing how the protocols achieve their goals and how they avoid pitfalls, the authors convey intricate design information that is essential for anyone venturing into the cryptographic protocol arena. Much of this analysis is probably not presented anywhere else, and certainly not in such a lucid fashion.

The last section unveils the secrets of several important standard protocols, including Kerberos, PEM, PGP, DCE, etc. What I liked best about this section was that it had details (message formats) combined with descriptions of the purpose of the protocol exchanges. It's rare to see something that serves as an intermediary reference in this fashion. One could use this for some reference functions in preference to programmer's manual. It's shorter and easier to understand.

Enjoyable

In addition to being enjoyable reading, this is the sort of book that gets used frequently by computer science students and professionals for reminders of how and why network security is achieved.

—Hilarie Orman
ho@earth.hpc.org

Call For Papers

IEEE Communications Magazine will publish a Special Issue on "Intranet Services and Communication Management" in October 1997.

Background

Intranets are corporate computer networks that use Internet software, hardware, middleware and protocol technologies to support computing and communication requirements within a company. Networks are being stressed to their limits as the corporations attempt to:

- Reduce their operating costs via automation,
- Continue to deploy more and more sophisticated co-operating multimedia capable desk-tops running latest versions of complex software
- Support multi-site decision making by giving access to private company information over internet and
- Provide sales and services using the Internet.

The costs for maintenance and keeping desk-tops up-to-date can be as high as \$10,000.00 per desk-top per year.

Although intranets facilitate (1) networked computing and network download of software upgrades, (2) server-centric soft servicing of clients and (3) use of quasi-sophisticated desk-tops running lightweight applications without compromising productivity, the real technological challenges which the businesses will be facing are: reliability, response time, auditability, authentication, privacy, security, etc.

Intranet is also providing new way to organise an Enterprise and new ways of doing business (through interaction between Intranets and public networks/Internet). Issues such as re-engineering of the enterprise gain a new dimension and partnering as well as virtual enterprises becomes possible in novel ways.

This feature topic/special issue will publish four to six carefully selected peer-reviewed articles discussing vision, technical and business aspects of intranet services and communication management by authors from industry and academia, practicing engineers and turn-key solution providers, and the standardization and management authorities from all over the world.

Submissions

Paper can be submitted to either of the guest editors listed below.

Important dates

Submission Deadline:	April 15, 1997
Acceptance Notification:	June 15, 1997
Revised Manuscript Due:	August 1, 1997
Publication:	October 1997

Guest Editors

Bhumip Khasnabish
GTE Labs. Inc., MS-48
40 Sylvan Road
Waltham, MA 02254
USA
Tel: +1-617-466-2080
Fax: +1-617-890-9320
E-Mail: bhumip@gte.com

Roberto Saracco
CSELT
via Reiss Romoli 270
Turin 10148
ITALY
+39-11-228-6906
+39-11-228-5685
roberto.saracco@cse.lt.stet.it

Future NetWorld+Interop Dates and Locations

NetWorld+Interop 97	Singapore	April 7–11, 1997
NetWorld+Interop 97	Las Vegas, NV	May 5–9, 1997
NetWorld+Interop 97	Frankfurt, Germany	May 12–15, 1997
NetWorld+Interop 97	Tokyo, Japan	June 2–6, 1997
NetWorld+Interop 97	Atlanta, GA	October 6–10, 1997
NetWorld+Interop 97	Paris, France	October 20–23, 1997
NetWorld+Interop 97	London, England	October 27–30, 1997
NetWorld+Interop 97	Sydney, Australia	November 25–28, 1997

All dates are subject to change.

More information

Call 1-800-INTEROP or +1-415-578-6900 for more information. Or send e-mail to info@interop.com or fax to +1-415-525-0194. For the latest information about Interop DotCom and NetWorld+Interop as well as other SOFTBANK produced events, check our *Interop Online* home page at <http://www.interop.com>. This is also where you will find *ConneXions Online* in the future.

NetWorld+Interop is produced by SOFTBANK Exposition and Conference Company, 303 Vintage Park Drive, Foster City, California 94404–1138, USA.

Write to *ConneXions* !

Even though this is the last printed issue of this journal, we would still love to hear your comments, suggestions and questions about anything you have read in *ConneXions*. Our editorial address is given below. Use it for letters to the Editor, requests for the index of back issues, questions about particular articles etc.:

ConneXions—The Interoperability Report

303 Vintage Park Drive

Suite 201

Foster City

California 94404–1138

USA

Phone: +1 415-578-6900 or 1-800-INTEROP (Toll-free in the USA)

Fax: +1 415-525-0194

E-mail: connexions@interop.com

URL: <http://www.interop.com>

Subscription information

For questions about your subscription please call our customer service hotline: 1-800-575-5717 or +1 610-892-1959 outside the USA. This is the number for our subscription agency, Seybold Publications. Their fax number is +1 610-565-1858. The mailing address for subscription payments is: P.O. Box 976, Media, PA 19063–0976.

This publication was distributed on an “as is” basis, without warranty. Neither the publisher nor any contributor shall have any liability to any person or entity with respect to any liability, loss, or damage caused or alleged to be caused, directly or indirectly, by the information contained in *ConneXions—The Interoperability Report*®

CONNEXIONS

303 Vintage Park Drive
Suite 201
Foster City, CA 94404-1138
Phone: 415-578-6900
FAX: 415-525-0194

ADDRESS CORRECTION
REQUESTED

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD

Dr. Vinton G. Cerf
Senior Vice President, MCI Telecommunications
President, The Internet Society (1992 – 1995)

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute



Printed on recycled paper

Subscribe to CONNEXIONS

U.S./Canada ☐ \$195. for 12 issues/year

All other countries ☐ \$245. for 12 issues/year

Name _____ Title _____

Company _____ E-mail _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

Fax () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card# _____ Exp.Date _____

Signature _____

Please return this application with payment to:

Back issues available upon request \$15./each
Volume discounts available upon request

CONNEXIONS

303 Vintage Park Drive, Suite 201
Foster City, CA 94404-1138
415-578-6900 FAX: 415-525-0194
connexions@interop.com

CONNEXIONS